Report of the CSAC Differential Privacy Working Group

May 25, 2021

Jay Breidt (convener), Deborah Balk, John Czajka,
Kathy Pettit, Allison Plyer (ex officio),
Kunal Talwar, Richelle Winkler, Joe Whitley

Differential Privacy as a technical achievement

- *WG previously noted that the Bureau's implementation of Differential Privacy (DP) at 2020 Census scale via TopDown Algorithm (TDA) is a **major technical achievement:***
  - largest application to date.
  - closely followed by other federal agencies.
  - closely watched by other countries, including Canada and UK.
- *WG applauds the Bureau's DP recognition by MIT's Technology Review as **one of 10 top Breakthrough Technologies of 2020** (along with anti-aging drugs, unhackable internet, and 7 others).*
  https://www.technologyreview.com/10-breakthrough-technologies/2020/#differential-privacy

I. Communications and engagement

- WG notes that it is essential that **pros and cons** of this innovative technology are…
  - studied and understood.
  - documented and disseminated.
- Serious concerns remain about how DP will affect the data's fitness for use in different use cases, so **communication and transparency** will be key to maintaining users' trust.
- *WG applauds the Bureau's ongoing efforts in communications:*
  - *blogs, webinars, newsletters and presentations about the forthcoming 2020 Census products and the evolving plans for applying DP.*
  - *commissioning of user handbooks to explain the implications of DP for different use cases and audiences.*

I. Communications and engagement, continued

- ***DRAFT RECOMMENDATIONS (for full CSAC consideration):***
- ***(I.a) The Bureau should continue its efforts to regularly communicate updates and engage various user groups during the decision-making process for setting the privacy-loss budget (PLB) and its allocation.***
- ***(I.b) The Bureau should actively engage with the stakeholder community - researchers, local and state government staff, other federal agency staff, and others - as the Bureau makes final decisions about the TopDown Algorithm to release the P.L. 94-171 Redistricting Data, and going forward with remaining products.***

I. Communications and engagement, continued:

- ***(I.c) The Bureau should communicate the factors used by the Data Stewardship Executive Policy (DSEP) Committee to set the PLB ("level of epsilon").***
- ***(I.d) The Bureau should publish 2020 Census data handbooks for data users targeted to different audiences (AIAN, federal agencies, data for rural areas, etc.) that parallel the handbooks created for the American Community Survey.***
- WG endorses the decision of the Bureau to publish plain-language explanations for users about practical implications of DP, starting with the P.L. 94-171 file and updated as new products are released.

II. PLB and custom geographies

- *WG commends the Bureau's efforts to improve the DP-adjusted estimates for off-spine geography and to increase the PLB allocation to block-level geography.*
- WG recognizes the importance of block-level data in constructing higher-level undefined geographies such as Congressional districts rather than their direct use.
- WG appreciates that higher-level custom geographies have better properties if they are closer to the geographic spine
- However, users continue to focus on accuracy at the block level.

II. PLB and custom geographies, continued:

- ***(II.a) DRAFT RECOMMENDATIONS (for full CSAC consideration): The Bureau should publish evaluations and examples to show how error declines with the aggregation of blocks into previously undefined geographies and to demonstrate that biases introduced by post-processing do not accumulate into larger errors for these undefined geographies, especially in low population density regions. The Bureau should produce these analyses as quickly as possible, as they will contribute critical evidence for making decisions about the needed PLB for the redistricting file.***

II. PLB and custom geographies, continued:

- The Bureau accepted our Fall 2020 recommendation that they should make “readily available tools for extrapolating from 2010 demonstration metrics to 2020 use cases.”
- WG understands that the Bureau is researching how to produce and provide these tools. Here, we add to the list of tools:
- ***(II.b) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should develop a set of tools to help users understand how error properties of custom geographies vary with distance from the geographic spine.***
- WG encourages the Bureau to practice user-centered design with the target audience in mind in developing these and other tools.
- WG requests an update on the progress of tools research and development, either to the DP working group or the full CSAC as appropriate.

II. PLB and custom geographies, continued:

- Metrics from April 28 demonstration data show that the reallocation of the PLB toward optimized block groups and blocks has:
  - improved the estimates (reduced error) for small off-spine geographies (such as incorporated places and minor civil divisions below 5,000 in population) and for census blocks.
  - increased tract-level errors substantially (4-fold when November 2020 is compared to April 28 with an overall PLB of 12.2), to the point where tracts seem out of line with other on-spine geographies.
- ***(II.c) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should consider whether this reduction in the quality of tract-level estimates represents an acceptable trade-off for the improvements achieved elsewhere or whether some reallocation of PLB to improve the quality of estimates for census tracts is merited.***

II. PLB and custom geographies, continued:

- Blocks with prison populations are especially important to identify clearly and count accurately for redistricting uses.
    - In some blocks, incarcerated people make up the entire population.
- Numbers of inmates are publicly published by other state and federal agencies.
- Potential inconsistencies between decennial census counts and numbers of inmates published elsewhere, possibly eroding data users' trust.
- The Bureau could draw on external numbers to improve decennial accuracy in these blocks.
- ***(II.d) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should investigate ways to keep prison block information as accurate as possible or at least consistent with data released by other sources, such as the Department of Justice. The Bureau should explore using externally published numbers of people in prisons in the post-processing in order to maintain consistency across data published from various sources, to increase data user trust, and to maintain accuracy without using the PLB or compromising privacy protection.***

III. Summary and use-case metrics

- Quality metrics at a finer scale are needed to help stakeholders, particularly in less populated places, understand the impact of DP on fitness of use for cases such as distributions of government funding or planning for community services.
- The previous Census Bureau response to the Fall 2020 recommendation - that users can create their own new metrics from privacy-protected microdata files - is not reasonable for the vast majority of users.
- ***(III.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should include more information about the range and pattern of error in releases of the Detailed Summary Metrics for future sprint cycles.***

III. Summary and use-case metrics, continued

- Split the current MALPE (Mean Algebraic Percent Error) statistic into the average negative relative error and the average positive relative error, rather than combining the two.
- Report the number of blocks in more detailed categories than the current ones of percent error greater than 5 percent (such as 5 to 10 percent, 10 to 20 percent, greater than 20 percent).
- Include the range from lowest to highest percent error.
- ***WG understands that this is additional work for limited Census Bureau staff time but believes these additions are needed for users to be able to evaluate how to appropriately use the privacy-protected data***.

III. Summary and use-case metrics, continued

- *WG appreciates the Bureau's continued inclusion of "impossible and improbable results" among its use case metrics.*
- April 28 metrics show that four of the eight impossible/improbable results have been reduced substantially (one to zero) by the combination of changes to the TDA and an increase in the PLB.
  - Blocks with children but no adults fell from 9.36% of the relevant universe to 1.47%
  - Substantial improvement, but will still bother users
- However, four of the results have changed little in frequency since the November 2020 release; for example:
  - 10.76% of census blocks with at least one occupied housing unit have a total household population of zero
  - 21.90% of all blocks have 100 percent occupancy in the DP estimates but not the published census data.

III. Summary and use-case metrics, continued

- WG is concerned that these impossible/improbable cases are especially problematic for users
- Detract from the Bureau's messaging about the impact of DP on the quality and usability of the data.
- ***(III.b) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should explicitly address these anomalies in its application of DP and its communication regarding such findings.***

IV. Helping users understand implications of post-processing

- *WG appreciates the Bureau's attention to improving its post-processing adjustments in order to reduce the error that these adjustments introduce into census data.*
- WG is concerned that post-processing error is still large
- Post-processing error may help to explain why the substantial increase in PLB from 4.5 to 12.2 did not produce larger reductions in the various error metrics (most were reduced by around one-half).
- ***(IV.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should communicate to users whatever success has been achieved in reducing post-processing error and provide evidence that post-processing error is not limiting in a substantial way the overall reduction in error achieved by the increase in the PLB.***

IV. Post-processing, continued

- The noisy counts prior to post-processing in the TDA give unbiased estimates, with analyzable error distributions.
- Having this data available would facilitate assessment of bias properties for the privacy-protected data, including potential positive biases created during post-processing, particularly in small domains due to nonnegativity constraints.
- A concern is that these small positive biases can accumulate as small domains are combined to create custom geographies.

IV. Post-processing, continued

- Therefore, WG suggests reasserting a Fall 2020 recommendation:
- ***(IV.b) DRAFT RECOMMENDATION (for full CSAC consideration): To facilitate assessment of bias properties for the privacy-protected data, the Bureau should release the non-post-processed measurements used in TDA, which are unbiased estimates with known error distributions. To address the Bureau's concerns that the release of such estimates would require extensive user guidance, the Bureau should consider releasing such data as a research product.***

V. Understanding and managing risks of reconstruction

- *WG appreciates the Bureau's recognition of the privacy risks posed by exact invariants.*
- Exact invariants impose hard constraints on the data, making reconstruction attacks both computationally easier and much more feasible, thus undermining the privacy technology.
- ***(V.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should continue to minimize the number of invariants they utilize.***

V. Risks of reconstruction, continued

- WG recognizes that Bureau's simulated reconstruction attack does not represent worst-case scenario
  - used only a subset of available tables.
  - stopped when the reconstruction risk was demonstrated.
- More sophisticated attacks will likely be able to go farther than the reconstruction attack reported by the Bureau.
- Even a partially successful reconstruction attack could have a lasting negative effect on participation, and not in a uniform way.

V. Risks of reconstruction, continued

- The Bureau may find it useful to measure the effectiveness of simulated reconstruction attacks in re-identifying individuals, **especially those who are different from those around them**, to enable better understanding of the privacy risks at a certain PLB.
- ***(V.b) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should publish more details about the reconstruction attack, including the distribution on demographics and within-block minority status of the confirmed re-identifications.***

V. Risks of reconstruction, continued

- WG does not have enough evidence yet to determine if a PLB of 12.2 will produce data accurate enough for redistricting.
- Implications of this PLB choice for other products down the road, particularly Demographics and Housing Characteristics (DHC), are unclear.
- PLB of 12.2 is quite high, and may not provide sufficient privacy.
- Actual risks of reconstruction have not been sufficiently quantified nor understood at a PLB of 12.2.
- ***(V.c) DRAFT RECOMMENDATION (for full CSAC consideration): The PLB should be selected to guard against attacks stronger than the attack reported by the Bureau.***
- WG notes that accuracy improvements to date are primarily due to algorithmic improvements.
- ***(V.d) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should continue to improve their algorithms such that the PLB can be minimized to achieve necessary levels of accuracy, with needed levels of protection.***

VI. Understanding implications of DP on different use cases

- The Bureau accepted Fall 2020 CSAC recommendation to publish further details on impacted geographic levels and variables (and their combination) and committed to conduct an assessment of the accuracy and trade-offs in future versions on an ongoing basis.
- WG recognizes the need for additional evaluation from a variety of data users and reiterates the importance of keeping these activities ongoing. Until substantially more analyses have been conducted, the risks of releasing 2020 Census data products to which DP has been applied (at various epsilon levels) are not known.

- ***(VI.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should continue to collect and study (internally or with collaborators) use cases across a wide range of uses, variables, geographies, and among a wide range of stakeholders both for the PL data and for subsequent data releases.***

VI. Different use cases, continued

- *WG commends the Bureau for its work to evaluate the redistricting use case by assessing the errors due to the TopDown Algorithm in the congressional districts created during post-2010 redistricting (Wright and Iramata, 2020)*
    - *Such work is highly useful for demonstrating how fit for use DP data are in the all-important redistricting use case.*
- ***(VI.b) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should extend this analysis to investigate how DP would have impacted congressional districts, and examples of smaller political districts such as precincts or school districts in low population areas that were drawn post 2010, using the latest demonstration data.***

VI. Different use cases, continued

- *WG applauds the Bureau's release of the 4/28/2021 demonstration data, which is allowing for additional evaluation from data users for a variety of uses.*
- Data released in the PL file are used for **many other applications beyond redistricting**, including federal/state/local government mandates and planning, that the Bureau has not, as far as the WG is aware, evaluated similarly to the redistricting use.
- Until substantially more analyses have been conducted, the **risks of releasing 2020 Census data products** to which DP has been applied (at various epsilon levels) are not known.
- WG reiterates the importance of **additional rigorous analysis** for different use cases:
  - analyses of impacts on funding formulas for federal agencies and congressional staffs.
  - analyses of impacts on legal mandates and regulatory practices, including protections for civil rights.

VI. Different use cases, continued

- Evaluation of racial and environmental justice impacts is a federally mandated example, in which **block level data on race/ethnic composition** are used to test for differential impacts of social and environmental goods and bads:
  - Environmental Protection Agency: evaluations of environmental justice.
  - US Forest Service: environmental impact assessments.
- **Executive Order on Advancing Racial Equity and Support for Underserved Communities Through the Federal Government (01/20/2021)**: calls extra attention to the importance of accurate data on small area race/ethnic composition for meeting these demands.

VI. Different use cases, continued

- ***(VI.c) DRAFT RECOMMENDATION (for full CSAC consideration): The Census Bureau should evaluate the impact of DP for racial equity uses, including Fair Housing Act and environmental justice, following the model of the redistricting evaluation (Wright and Iramata 2020).***

VII. DP and substate quality metrics

- *WG commends the Census Bureau for releasing a large number of quality metrics at the state level.*
- Quality metrics released for 2020 Census state level data have been extremely helpful in understanding the quality of that data.
    - For example, these metrics reveal that in Louisiana 0.91% of all addresses were resolved through true Count Imputation (a rate 4 times the national average of .23%).
- Releasing such metrics at the census tract level will help local planners understand the reliability of local area data, and where they may want to augment 2020 Census data with local administrative data for emergency response, road planning, and more.

VII. DP and substate quality metrics, continued

- Working group cannot currently envision a justification for applying DP to census tract level metrics such as:
  - number of housing units with counts imputed.
  - number of housing units enumerated by administrative data.
  - number of housing units enumerated by proxy.
  - number of NRFU housing units that were enumerated as nonexistent.
- Applying DP to the quality metrics will make them largely irrelevant, and will take part of the PLB away from important data products.

VII. DP and substate quality metrics, continued

- *(VII.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should not apply DP to substate quality metrics that are being released to the public -- either via the ASA task force or directly from the Bureau.*
- ***(VII.b) If the Bureau concludes that the quality metrics cannot be released without applying DP, we request that the Bureau specify how these metrics could be used in a reconstruction scenario, to justify their decision.***

VIII. Trade-offs of block/block group data and DHC

- WG recognizes that the risk of disclosure is greatest for block level data cross-tabulated by detailed characteristics, such as age and race/ethnicity, that will be part of the Demographics and Housing Characteristics (DHC) file.
- ***(VIII.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should examine the utility of and explore tradeoffs of not releasing all the detailed tables in the DHC file at the block and potentially block group levels, if this would allow for more accuracy in other geographic levels.***
- The goal would be to ensure that publicly released data are robust with a very clear indication of fitness for use.
- As the Bureau considers this strategy, they should engage with the user community to gain stronger fitness of use for other use cases.

IX. FSRDCs

- Federal Statistical Research Data Centers (FSRDCs) are valuable resources built up over the last two decades, developed as an early answer to disclosure avoidance. The ability to access data in a controlled, restricted environment is important.
- ***(IX.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should work on a plan to maintain FSRDC utility and access to high quality data.***
- FSRDCs would be especially important if the Bureau decides to restrict release of detailed block level data. Retaining this enclave approach, where researchers can access data where accuracy is not compromised by the differential privacy, would allow for the continuation of critical use cases.
- Research and conclusions stemming from FSRDC-based analyses would need to be assessed for privacy loss, with a portion of the PLB set aside for such uses.

X. Timeline for 2020 Census product releases

- Bureau's implementation of DP has followed an ambitious timeline.
- Many implications of DP implementation are not yet fully understood.
  - risk of reconstruction attacks based on different levels of the PLB has not been fully quantified.
  - fitness for use of legal and regulatory uses of the data have not been examined in full.
  - there are unquantified risks of failing to produce sufficiently accurate data for some legal and regulatory uses.
- *WG commends the Bureau for prioritizing research over speed of release in preparing the redistricting data. This is an important precedent for reducing risks (both in privacy loss and also in lack of fitness for use) before releasing data products.*

X. Timeline, continued

- ***(X.a) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should de-prioritize speed and prioritize performing the research necessary to understand and reduce the risk associated with privacy loss and insufficient fitness for use.***
- While CSAC has provided information about some key use cases, many critical use cases are still unknown, and CSAC anticipates risks associated with releasing data products that are not sufficiently accurate for these use cases.
- A more complete use case catalog, as CSAC previously recommended, is still needed to identify and mitigate such risks.

X. Timeline, continued

- ***(X.b) DRAFT RECOMMENDATION (for full CSAC consideration): The Bureau should delay additional releases after the September redistricting file to allow sufficient time for developing the required new algorithms, testing the implications of alternative allocations of the PLB, assessing the risk of privacy loss from various epsilons, assessing risks of releasing data that is not fit-for-use (particularly for legal applications of decennial census data products), and developing demonstration products to inform users of the likely accuracy of the data.***
- WG recognizes that data users will be inconvenienced by further delays in releases of decennial data products
- But such delays will likely increase the accuracy of the resulting products while improving privacy protection, as the Bureau's techniques for developing and deploying DP are rapidly evolving.
- In addition, taking the time needed to do this work well will yield downstream benefits for other federal statistical agencies.

Discussion